

# A Distance Measure for Attention Focusing and Anomaly Detection in Systems Monitoring

Richard J. Doyle

Artificial Intelligence Group  
Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, CA 91109-8099  
rdoyle@aig.jpl.nasa.gov

## Abstract

Any attempt to introduce automation into the monitoring of complex physical systems must start from a robust anomaly detection capability. This task is far from straightforward, for a single definition of what constitutes an anomaly is difficult to come by. In addition, to make the monitoring process efficient, and to avoid the potential for information overload on human operators, attention focusing must also be addressed. When an anomaly occurs, more often than not several sensors are affected, and the partially redundant information they provide can be confusing, particularly in a crisis situation where a response is needed quickly.

Previous results on extending traditional anomaly detection techniques are summarized. The focus of this paper is a new technique for attention focusing. The technique involves reasoning about the distance between two frequency distributions, and is used to detect both anomalous system parameters and "broken" causal dependencies. These two forms of information together isolate the locus of anomalous behavior in the system being monitored.

## 1 Introduction

Mission Operations personnel at NASA have the task of determining, from moment to moment, whether a space platform is exhibiting behavior which is in any way anomalous, which could disrupt the operation of the platform, and in the worst case, could represent a loss of ability to achieve mission goals. A traditional technique for assisting mission operators in space platform health analysis is the establishment of alarm thresholds for sensors, typically indexed by operating mode, which summarize which ranges of sensor values imply the existence of anomalies. Another established technique for anomaly detection is the comparison of predicted values from a simulation to actual values received in telemetry. However, experienced mission operators reason about more than alarm threshold crossings and discrepancies between predicted and actual to detect anomalies: they may ask whether a sensor is behaving differently than it has in the past, or whether a current behavior may lead to—the particular bane of operators—a rapidly developing alarm sequence.

Our approach to introducing automation into real-time systems monitoring is based on two observations: 1) mission operators employ multiple methods for recognizing anomalies, and 2) mission operators do not and should not interpret all sensor data all of the time. We seek an approach for determining from moment to moment which of the available sensor data is most informative about the presence of anomalies occurring within a system. We term this process *sensor selection* and we have implemented a prototype selective monitoring system called SELMON [6; 7].

The SELMON system has its origins in a sensor planning system called GRIPE [5] which planned information gathering activities to verify the execution of robot task plans. Other model-based monitoring systems include Dvorak's MIMIC, which performs robust discrepancy detection for continuous dynamic systems [8; 9], and DeCoste's DATMI, which infers system states from incomplete sensor data [4]. The SELMON work complements other work within NASA on empirical and model-based methods for fault diagnosis of aerospace platforms [1; 10; 11; 13].

## 2 Background: The SELMON Approach

How does a human operator or a machine observing a complex physical system decide when something is going wrong? Abnormal behavior is always defined as some kind of departure from normal behavior. Unfortunately, there appears to be no single, crisp definition of "normal" behavior. In the traditional monitoring technique of limit sensing, normal behavior is predefined by nominal value ranges for sensors. A fundamental limitation of this approach is the lack of sensitivity to context. In the other traditional monitoring technique of discrepancy detection, normal behavior is obtained by simulating a model of the system being monitored. This approach, while avoiding the insensitivity to context of the limit sensing approach, has its own limitations. The approach is only as good as the system model. In addition, normal system behavior typically changes with time, and the model must continue to evolve. Given these limitations, it can be difficult to distinguish genuine anomalies from errors in the model.

Noting the limitations of the existing monitoring techniques, we have developed an approach to monitoring which is designed to make the anomaly detection process more robust, to reduce the number of undetected anomalies (false negatives). Towards this end, we introduce multiple anomaly

models, each employing a different notion of "normal" behavior.

## 2.1 Empirical Anomaly Detection Methods

In this section, we briefly describe the empirical methods that we use to determine, from a local viewpoint, when a sensor is reporting anomalous behavior. These measures use knowledge about each individual sensor, without knowledge of any relations among sensors.

### Surprise

An appealing way to assess whether current behavior is anomalous or not is via comparison to past behavior. This is the essence of the *surprise* measure. It is designed to highlight sensor which behaves other than it has historically. Specifically, *surprise* uses the historical frequency distribution for the sensor in two ways: To determine the likelihood of the given current value of the sensor (*unusualness*), and to examine the relative likelihoods of different values of the sensor (*informativeness*). It is those sensors which display unlikely values when other values of the sensor are more likely which get a high *surprise* score. *Surprise* is not high if the only reason a sensor's value is unlikely is that there are many possible values for the sensor, all equally unlikely.

### Alarm

Alarm thresholds for sensors, indexed by operating mode, typically are established through an off-line analysis of system design. The notion of *alarm* in SELMON extends the usual one bit of information (the sensor is in alarm or it is not), and also reports how much of the alarm range has been traversed. Thus a sensor which has gone deep into alarm gets a higher score than one which has just crossed over the alarm threshold.

### Alarm Anticipation

The *alarm anticipation* measure in SELMON performs a simple form of trend analysis to decide whether or not a sensor is expected to be in alarm in the future. A straightforward curve fit is used to project when the sensor will next cross an alarm threshold, in either direction. A high score means the sensor will soon enter alarm or will remain there. A low score means the sensor will remain in the nominal range or emerge from alarm soon.

### Value Change

A change in the value of a sensor may be indicative of an anomaly. In order to better assess such an event, the *value change* measure in SELMON compares a given value change to 10 historical value changes seen on that sensor. The score reported is based on the proportion of previous value changes which were less than the given value change. It is maximum when the given value change is the greatest value change seen to date on that sensor. It is minimum when no value change has occurred in that sensor.

## 2.2 Model-Based Anomaly Detection Methods

Although many anomalies can be detected by applying anomaly models to the behavior reported at individual sensors, robust monitoring also requires reasoning about interactions occurring in a system and detecting anomalies in behavior reported by several sensors.

### Deviation

The *deviation* measure is our extension of the traditional method of discrepancy detection. As in discrepancy detection, comparisons are made between predicted and actual sensor values, and differences are interpreted to be indications of anomalies. This raw discrepancy is entered into a normalization process identical to that used for the *value change* score, and it is this representation of relative discrepancy which is reported. The *deviation* score for a sensor is minimum if there is no discrepancy and maximum if the discrepancy between predicted and actual is the greatest seen to date on that sensor.

*Deviation* only requires that a simulator be available in any form for generating sensor value predictions. However, the remaining *sensitivity* and *cascading alarms* measures require the ability to simulate and reason with a causal model of the system being monitored.

### Sensitivity and Cascading Alarms

*Sensitivity* measures the potential for a large global perturbation to develop from current state. *Cascading alarms* measures the potential for an alarm sequence to develop from current state. Both of these anomaly measures use an event-driven causal simulator [2; 12] to generate predictions about future states of the system, given current state. Current state is taken to be defined by both the current values of system parameters (not all of which may be sensed) and the pending events already resident on the simulator agenda. The measures assign scores to individual sensors according to how the system parameter corresponding to a sensor participates in, or influences, the predicted global behavior. A sensor will have its highest *sensitivity* score when behavior originating at that sensor causes all sensors causally downstream to exhibit their maximum value change to date. A sensor will have its highest *cascading alarms* score when behavior originating at that sensor causes all sensors causally downstream to go into an alarm state.

## 2.3 Previous Results

In order to assess whether SELMON increased the robustness of the anomaly detection process, we performed the following experiment: We compared SELMON performance to the performance of the traditional limit sensing technique in selecting critical sensor subsets specified by a Space Station Environmental Control and Life Support System (ECLSS) domain expert, sensors seen by that expert as useful in understanding episodes of anomalous behavior in actual historical data from ECLSS testbed operations.

The experiment asked the following specific question: How often did SELMON place a "critical" sensor in the top half of its sensor ordering based on the anomaly detection measures?

The performance of a random sensor selection algorithm would be expected to be about 50%; any particular sensor would appear in the top half of the sensor ordering about half the time. Limit sensing detected the anomalies 76.3% of the time. SELMON detected the anomalies 95.1% of the time.

These results show SELMON performing considerably better than the traditional practice of limit sensing. They lend credibility to our premise that the most effective monitoring system is one which incorporates several models of anomalous behavior. Our aim is to offer a more complete, robust set of techniques for anomaly detection to make human oper-

ators more effective, or to provide the basis for an automated monitoring capability.

The following is a specific example of the value added of SELMON. During an episode in which the ECLSS pre-heater failed, system pressure (which normally oscillates within a known range) became stable. This “abnormally normal” behavior is not detected by traditional monitoring methods because the system pressure remains firmly in the nominal range where limit sensing fails to trigger. Furthermore, the fluctuating behavior of the sensor is not modeled; the predicted value is an averaged stable value which fails to trigger discrepancy detection. See [7; 6] for more details on these previous results in evaluating the SELMON approach.

### 3 Attention Focusing

A robust anomaly detection capability provides the core for monitoring, but only when this capability is combined with attention focusing does monitoring become both robust and efficient. Otherwise, the potential problems of information overload and too many false positives may defeat the utility of the monitoring system.

The attention focusing technique developed here uses two sources of information: historical data describing nominal system behavior, and causal information describing which pairs of sensors are constrained to be correlated, due to the presence of a dependency. The intuition is that the origin and extent of an anomaly can be determined if the misbehaving system parameters and the misbehaving causal dependencies can be determined. Such information also supports reasoning to distinguish whether sensors, system parameters or mechanisms are misbehaving due to the fact that the signature of “broken” nodes and arcs in the causal graph are distinguishable. See Figure 1.

For example, the expected signature of an anomalous sensor includes the node of the sensor itself and the immediately adjacent arcs corresponding to the causal dependencies that the sensor participates in directly. The intuition is that the actual system is behaving normally so the locus of “brokenness” is isolated to the sensor and the set of adjacent causal dependencies which attempt to reconcile the bogus value reported by the sensor.

The expected signature of an anomalous system parameter also includes nodes and arcs which are downstream in the causal graph from the node corresponding to the system parameter. The intuition here is that the misbehavior, being in the actual system, will propagate.

The expected signature of an anomalous mechanism also includes arcs and nodes causally downstream from the arc corresponding to the mechanism. Once again, the intuition is that the misbehavior is in the system itself, and it will propagate. The way to distinguish this case from the anomalous system parameter case is to examine all input arcs (assuming there are more than one) to the most causally prior node in the “broken” subgraph.

#### 3.1 Two Additional Measures

While SELMON runs, it computes incremental frequency distributions for all sensors being monitored. These frequency distributions can be saved as a method for capturing behavior from any episode of interest. Of particular interest are




Figure 1: Anomalous System Parameters, Sensors and Mechanisms.

historical distributions which correspond to nominal system behavior.

To identify an anomalous sensor, we apply a distance measure, defined below, to the frequency distribution which represents recent behavior to the historical frequency distribution representing nominal behavior. We call the measure simply *distance*. To identify a “broken” causal dependency, we first apply the same distance measure to the historical frequency distributions for the cause sensor and the effect sensor. This reference distance is a weak representation of the correlation that exists between the values of the two sensors due to the causal dependency. This reference distance is then compared to the distance between the frequency distributions based on recent data of the same cause sensor and effect sensor. The difference between the reference distance and the recent distance is the measure of the “brokenness” of the causal dependency. We call this measure *causal distance*.

#### 3.2 Desired Properties of the Distance Measure

Define a distribution  $D$  as the vector  $d_i$  such that

$$\forall i, 0 \leq d_i \leq 1$$

and

$$\sum_{i=0}^{n-1} d_i = 1$$

For a sensor  $S$ , we assume that the range of values for the sensor has been partitioned into  $n$  contiguous subranges which exhaust this range. We construct a frequency distribution as a

vector  $D_s$  of length  $n$ , where the value of  $d_i$  is the frequency with which  $S$  has displayed a value in the  $i$ th subrange.

If our aim was only to compare different frequency distributions of the same sensor, we could use a distance measure which required the number of partitions, or bins in the two distributions to be equal, and the range of values covered by the distributions to be the same. However, since our aim is to be able to compare the frequency distributions of different sensors, these conditions must be relaxed.

Before defining the other desired properties of the distance measure, we define two special types of frequency distribution. Let  $F$  be the random, or flat distribution where  $\forall i, d_i = \frac{1}{n}$ . Let  $S_i$  be the Set of "spike" distributions where  $d_i = 1$  and  $\forall j \neq i, d_j = 0$ .

We seek a distance measure for frequency distributions with the following properties:

#### Distance

$$\forall D_1, D_2, \Delta(D_1, D_2) \geq 0$$

This properly merely defines the measure as a distance measure.

#### Identity

$$\forall D, \Delta(D, D) = 0$$

#### Symmetry

$$\forall D_1, D_2, \Delta(D_1, D_2) = \Delta(D_2, D_1)$$

We do not wish to emphasize whether we are comparing recent data to historical data or vice versa, or cause data to effect data or vice versa.

#### Spike Distinctness

$$\forall i \neq j, \Delta(S_i, S_j) > 0$$

We wish the set of  $S_i$  to be distinguishable,

#### Spike Ordering

$$\forall i, \Delta(S_i, S_{i+1}) < \Delta(S_i, S_{i+2})$$

The distance measure should preserve the fact that there is an ordering on the bins.

#### Spike Equidistance

$$\forall i \neq j, \Delta(S_i, S_{i+1}) = \Delta(S_j, S_{j+1})$$

There should be no difference in weighting of the spike distributions.

#### Spike/Flat Equidistance

$$\forall i \neq j, \Delta(S_i, F) = \Delta(S_j, F)$$

The difference between any spike distribution and the flat distribution is to be the same.

$$\text{Extrema } \forall D_1, D_2 \forall i, \Delta(D_1, D_2) \leq \Delta(S_i, F)$$

Any spike distribution and the flat distribution arc to be considered the most different. All other distributions fall in between.

### 3.3 The Distance Measure

The distance measure is computed by projecting the two distributions into the two-dimensional space  $[f, s]$  in polar coordinates and taking the euclidian distance between the projections.

Define the "flatness" component  $f(D)$  of a distribution as follows:

$$\sum_{i=0}^{n-1} \frac{1}{2} \left| \frac{1}{n} - d_i \right|$$




Figure 2: The function  $A(D, D_2)$ .

This is simply the sum of the bin-by-bin differences between the given distribution and  $F$ . Note that  $0 \leq f(D) \leq 1$ . Also,  $f(S_i) \rightarrow 1$  as  $n \rightarrow \infty$ .

Define the "spikeness" component  $s(D)$  of a distribution as:

$$\sum_{i=0}^{n-1} \phi \frac{i}{n-1} d_i$$

This is simply the centroid value calculation for the distribution. The weighting factor  $\phi$  will be explained in a moment. Once again,  $0 \leq s(D) \leq 1$ .

Now take  $[f, s]$  to be polar coordinates  $[r, t]$ . This maps  $F$  to the origin and the  $S_i$  to points along an arc on the unit circle. See Figure 2.

By inspection, the *Spike Distinctness*, *Spike Ordering* and *Spike/Flat Equidistance* properties are satisfied. The *Spike Equidistance* property is satisfied because there is no unequal weighting applied in the centroid calculation. The *Distance*, *Identity* and *Symmetry* properties follow from taking the euclidian distance between the projections of the distributions. The *Extrema* property is satisfied by taking  $\phi = \frac{\pi}{3}$ . This choice of  $\phi$  guarantees that  $\Delta(S_0, S_{n-1}) = \Delta(F, SO) = A(F, S_{n-1}) = 1$  and all other distances in the region which is the range of  $A$  arc by inspection  $< 1$ .

*Insensitivity* to the number of bins in the two distributions and the range of values encoded in the distributions is provided by the  $[f, s]$  projection function, which abstracts away from these properties of the distributions.

We may note in passing that the distance measure described here may be easily modified to apply to continuous distributions, when theoretical models of the behavior of a system are available. The centroid calculation of this component is easily accomplished, and the  $f$  component involves merely the integral of a difference, which may be accomplished numerically if necessary.

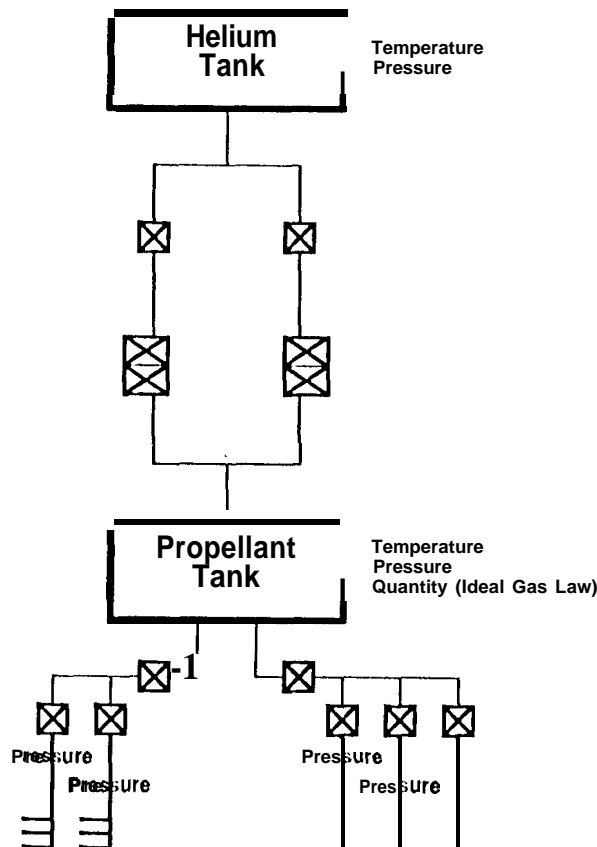


Figure 3: The Forward Reactive Control System (FRCS) of the Space Shuttle.

### 3.4 Results

In this section, we report on the results of applying the distribution distance measure to the task of focusing attention in monitoring. The distribution distance measure is used to identify misbehaving nodes (*distance*) and arcs (causal *distance*) in the causal graph of the system being monitored, or equivalently, detect and isolate the extent of anomalies in the system being monitored.

#### 3.4.1 A Space Shuttle Propulsion Subsystem

Figure 3 and Figure 4, show respectively, a schematic diagram and a causal graph for a portion of the Forward Reactive Control System (FRCS) of the Space Shuttle. A full causal graph for the Reactive Control System, comprising the Forward, Left and Right RCS, was developed with the domain expert.

#### 3.4.2 Attention Focusing Examples

SELMON was run on seven episodes describing nominal behavior of the FRCS. The frequency distributions collected during these runs were merged. Reference distances were computed for sensors participating in causal dependencies.

SELMON was then run on 13 different fault episodes, representing faults such as leaks, sensor failures and regulator failures. Two of these episodes will be examined here; results were similar for all episodes. In each fault episode, and for each sensor, the distribution distance measure was

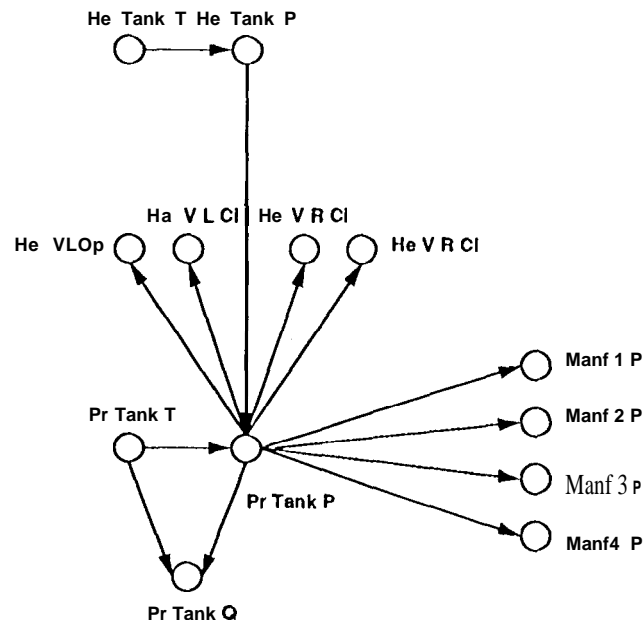


Figure 4: Causal Graph for the FRCS.

applied to the incremental frequency distribution collected during the episode and the historical frequency distribution from the merged nominal episodes. These distances were a measure of the "brokenness" of nodes in the causal graph; i.e., instantiations of the *distance* measure.

New distances were computed between the distributions corresponding to sensors participating in causal dependencies. The differences between the new distances and the reference distances for the dependencies were a measure of the "brokenness" of arcs in the causal graph; i.e., instantiation of the *causal distance* measure.

The first episode involves a leak affecting the first and second manifolds (jets) on the oxidizer side of the FRCS. The pressures at these two manifolds drop to vapor pressure. The dependency between these pressures and the pressure in the propellant tank is severed because the valve between the propellant tank and the manifolds is closed. Thus there are two anomalous system parameters (the manifold pressures) and two anomalous mechanisms (the agreement between the propellant and manifold pressures when the valve is open).

The *distance* and *causal distance* measures computed for nodes and arcs in the FRCS causal graph reflect this faulty behavior. See Figure 5. (To visualize how the distribution distance measure circumscribes the extent of anomalies, the coloring of nodes and the width of arcs in the figure are correlated with the magnitudes of the associated *distance* and *causal distance* scores).

The second episode involves an overpressurization of the propellant tank due to a regulator failure. Onboard software automatically attempts to close the valves which isolate the helium tank from the propellant tank. One of the valves sticks and remains open.

The *distance* and *causal distance* measures isolate both the misbehaving system parameters (propellant pressure and valve status indicators) and the misbehaving mechanisms (be-




Figure 5: A leak fault.




Figure 6: A pressure regulator fault.

tween the helium and propellant tank pressure and between the propellant tank pressure and the valve status indicators). Overpressurization of the propellant tank also alters the usual relation between propellant tank pressure and manifold pressures. See Figure 6.

## 4 Discussion

The *distance* and *causal distance* measures bawd on the **distribution** distance measure combine two concepts: 1) **empirical** data alone can bc an **effective** model of behavior, and 2) the **existence** of a causal **dependency** between two **parameters** implies **that** their values are somehow correlated. The *causal distance* measure **constructs** a model of the **correlation** between two causal] y related parameters, capturing the general notion of constraint in an admittedly **abstract** manner. Nonetheless, these models of constraint arising from **causality** provide surprising discriminatory power for **determining** which causal dependencies (and corresponding system mechanisms) are misbehaving. (In the *distance* measure for detecting misbehaving system parameters, we are simply using the degenerate constraint of expected equality between historical and recent behavior.)

## 4.1 Monitoring Architecture

The attention focusing capability provided by the *distance* and *causal distance* measures can be combined with the multiple-viewpoint anomaly detection capability already developed in SELMON to construct a general monitoring architecture.

The multiple anomaly measures (including the *distance* and *causal distance* measures, which are anomaly detection measures in their own right) provide continuous anomaly detection capability. All of these measures are normalized to the range  $[0,1]$  so their sensitivity, individually or collectively, can be fine-tuned for the behavior of particular monitored systems. Whenever a detected anomaly is announced, the extent of the

anomaly is isolated by applying the results of the *distance* and *causal distance* measures to the causal graph of the **system**. If **SEIMON** is supporting a human operator, the operator's **attention** is focused on the locus of the anomaly, rather than the potentially long and confusing list of the individual **manifestations** of the anomaly.

Once an anomaly is **detected** and circumscribed, the **operator** can once again use the set of anomaly **detection** measures to **gather** supplementary, multiple-viewpoint information about the **detected** anomaly at multiple probe points.

## 4.2 Anomaly Characterization

Most model-based reasoning work has focused on diagnosis, treating monitoring as a “front-end”, with discrepancy detection usually chosen as the monitoring technique. The SELMON work suggests modifications to this view.

Monitoring is a complex, subtle and important task in its own right. The most sophisticated diagnosis engine is of **limited utility** if it is unreliably invoked by a weak anomaly detection module.

The monitoring/diagnosis distinction actually defines two poles of a continuum [3]. At one end is **anomaly detection**. The goal of anomaly detection is simply to **determine** if an anomaly exists. General models of what constitutes an anomaly are **utilized**, with **limited reference** to explicit **behavior** models. Reasoning is local rather than global.

Next in the continuum is anomaly characterization. The goal here is to describe the extent of anomalous behavior. Again, the use of explicit behavior models is limited, but reasoning now encompasses a global view of the system. The empirical anomaly detection and attention focusing capabilities of SELMON correspond to anomaly detection and anomaly characterization as defined here. The model-based anomaly detection capabilities of SELMON also are examples of anomaly characterization.

Next comes fault isolation. Reasoning now is **refined** from

- anomaly extent to anomaly source. Explicit behavior models may be used, but not explicit fault models.

Finally comes full-fledged fault diagnosis, which includes an explanation of how the proposed fault produced the anomalous behavior. Explicit fault models may be referenced to verify hypotheses.

In actual real-time monitoring practice, operators perform anomaly detection and characterization routinely, and fault isolation when enough information is available to support their reasoning. Fault diagnosis is typically done off-line.

## 5 Future Work

Several issues need to be examined to continue the evaluation of the attention focusing technique based on the distribution distance measure and its utility in monitoring.

We need to understand the sensitivity of the technique to how sensor value ranges are partitioned. Clearly the discriminatory power of the distribution distance measure is related to the resolution provided by the number of bins and the bin boundaries. The results reported here are encouraging for the number of FRCS sensor bins were in many cases as low as three and in no cases more than eight.

We also need to understand the suitability of the technique for systems which have many modes or configurations. We would expect that the discriminatory power of the technique would be compromised if the distributions describing behaviors from different modes were merged. Thus the technique requires that historical data representing nominal behavior is separable for each mode. If there are many modes, at the very least there is a data management task. A capability for tracking mode transitions is also required. An unsupervised learning system which can enumerate system modes from historical data and enable automated classification would solve this problem nicely.

The discriminatory power of the causal distance measure might be enhanced by retaining the flatness/spikeness distinction. For many linear functions, different input distributions may map to value-shifted but similarly shaped output distributions. In other words, the spikeness component may vary while the flatness component may be relatively invariant. It may be possible to distinguish the case where misbehavior is the result of bogus values being propagated through still correctly functioning mechanisms.

It should be possible to describe the temporal (along with the causal/spatial) extent of anomalies by incrementally comparing recent sensor frequency distributions calculated from a "moving window" of constant length with static reference frequency distributions.

## 6 Towards Applications

The approach described in this paper has usability advantages over other forms of model-based reasoning. The overhead involved in constructing the causal and behavioral model of the system is minimal. The behavioral model is derived directly from actual data; no offline modeling is required. The causal model is of the simplest form, describing only the existence of dependencies. For the Shuttle RCS, a 198-node, 196-arc causal graph was constructed in a single one and one half hour session between the author and the domain expert.

SELMON is being applied at the NASA Johnson Space Center as a monitoring tool for Space Shuttle Operations and Space Station Operations. Current application efforts include the one for the Frepulsion (PROP) flight control discipline reported on here, and one for the Thermal (EECOM) flight control discipline. EECOM wishes in particular to be able to know and reason about how actual orbiter thermal performance differs from predictions generated by an available mathematical model of orbiter thermal performance. An operational SELMON prototype, available starting with the recent Hubble Repair mission is available for evaluation by all flight control disciplines, only requiring that a list of sensors "owned" by that discipline be provided.

At the Jet Propulsion Laboratory, we are looking at the problem of onboard downlink determination for the Pluto Fast Flyby project, now in its early planning phase. The spacecraft will have limited communications bandwidth and it will not be possible to transmit all onboard-collected sensor data. Only eight hours of coverage from the Deep Space Network will be available per week. The challenge is to devise a method for constructing a suitable summary of a week's worth of sensor data guaranteed to report on any anomalies which occurred. The anomaly detection and attention focusing capabilities of SELMON may be well-matched to this task.

## 7 Summary

We have described the properties and performance of a distance measure used to identify misbehavior at sensor locations and across mechanisms in a system being monitored. The technique enables the locus of an anomaly to be determined. This attention focusing capability is combined with a previously reported anomaly detection capability in a robust, efficient and informative monitoring system called SELMON. SELMON is being applied in mission operations at NASA.

## 8 Acknowledgements

The members of the SELMON team are Len Charest, Harry Porta, Nicolas Rouquette and Jay Wyatt. Dan Clancy has also worked on the SELMON project recently. Usama Fayyad contributed work on the empirical anomaly detection measures. Steve Chien contributed work on the model-based anomaly detection measures. Matt Barry and Charles Robertson provided valuable discussion during the development of the distance measure. Matt Barry also served invaluable as domain expert for the Shuttle FRCS.

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

## References

- [11] K. H. Abbott, "Robust Operative Diagnosis as Problem Solving in a Hypothesis Space," *7th National Conference on Artificial Intelligence*, St. Paul, Minnesota, August 1988,
- [21] L. Charest, Jr., N. Rouquette, R. Doyle, C. Robertson, and J. Wyatt, "MESA: An Interactive Modeling and Simulation Environment for Intelligent Systems Automa-

tion," *Hawaii International Conference on System Sciences*, Maui, Hawaii, January 1994.

- [31] D. Clancy, "Anomaly Characterization," *Internal Report*, Jet Propulsion Laboratory, September 1993.
- [41] D. DeCoste, "Dynamic Across-Time Measurement Interpretation," *8th National Conference on Artificial Intelligence*, Boston, Massachusetts, August 1990.
- [51] R. J. Doyle, D. J. Atkinson, and R. S. Doshi, "Generating Perception Requests and Expectations to Verify the Execution of Plans," *5th National Conference on Artificial Intelligence*, Philadelphia, Pennsylvania, August 1986.
- [61] R. J. Doyle, L. Charest, Jr., N. Rouquette, J. Wyatt, and C. Robertson, "Causal Modeling and Event-driven Simulation for Monitoring of Continuous Systems," *Computers in Aerospace 9*, American Institute of Aeronautics and Astronautics, San Diego, California, October 1993.
- [71] R. J. Doyle, S. A. Chien, U. M. Fayyad, and E. J. Wyatt, "Focused Real-time Systems Monitoring Based on Multiple Anomaly Models," *7th International Workshop on Qualitative Reasoning*, Eastsound, Washington, May 1993.
- [81] D. L. Dvorak and B. J. Kuipers, "Model-Based Monitoring of Dynamic Systems: 11th International Conference on Artificial Intelligence," Detroit, Michigan, August 1989.
- [91] D. L. Dvorak, "Monitoring and Diagnosis of Continuous Dynamic Systems using Semiquantitative Simulation," Ph.D. dissertation, The University of Texas at Austin, May 1992.
- [101] T. Hill, W. Morris, and C. Robertson, "Implementing a Real-time Reasoning System for Rebus{ Diagnosis," *6th Annual Workshop on Space Operations Applications and Research*, Houston, Texas, August 1992.
- [111] J. Muratorc, T. Heindel, T. Murphy, A. Rasmussen, and R. McFarland, "Space Shuttle Telemetry Monitoring by Expert Systems in Mission Control," in *Innovative Applications of Artificial Intelligence*, H. Schorr and A. Rappaport (eds.), AAAI Press, 1989.
- [121] N. F. Rouquette, S. A. Chien, and L. K. Charest, Jr., "Event-driven Simulation in SILMON: An Overview of EDSF," JPL Publication 92-23, August 1992.
- [131] E. A. Scarf, J. R. Jamieson, and E. New, "Deriving Fault Location and Control from a Functional Model," *3rd IEEE Symposium on Intelligent Control*, Arlington, Virginia, 1988.